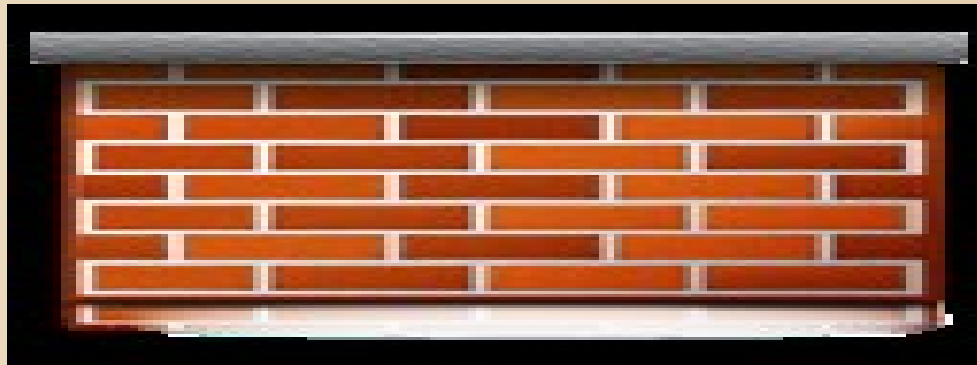


FIREWALLS



MARTIN NJAU
Systems Administrator
mnjau@kenet.or.ke

28th October 2015



kenet
Kenya Education Network

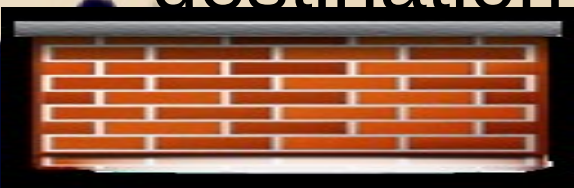
Introduction

- A firewall is a device or set of devices designed to permit or deny network transmissions based on a set of rules and is frequently used to protect networks from unauthorized access while permitting legitimate communications to pass.
- Many personal computer operating systems include software-based firewalls to protect against threats from the public Internet.
- Examples of firewalls include OpenBSD PF, IPFILTER, IPFW and iptables.



Introduction Cont'd

- Firewalls make it possible to filter incoming and outgoing traffic that flows through your system.
- A firewall can use one or more sets of “rules” to inspect the network packets as they come in or go out of your network connections and either allows the traffic through or blocks it.
- The rules of a firewall can inspect one or more characteristics of the packets, including but not limited to the protocol type, the source or destination host address, and the source or destination port.



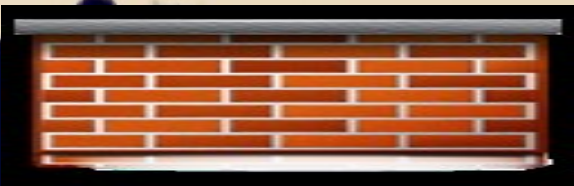
Functions of Firewalls

- Firewalls can greatly enhance the security of a host or a network. They can be used to do one or more of the following things:
 - To protect and insulate the applications, services and machines of your internal network from unwanted traffic coming in from the public Internet.
 - To limit or disable access from hosts of the internal network to services of the public Internet.
 - To support network address translation (NAT), allows your internal network to use private addresses and share a single



Firewall Concepts

- There are two basic ways to create firewall rulesets:
 - An exclusive firewall allows all traffic through except for the traffic matching the ruleset.
 - An inclusive firewall does the reverse. It only allows traffic matching the rules through and blocks everything else.
- An inclusive firewall offers much better control of the outgoing traffic, making it a better choice for systems that offer services to the public Internet.



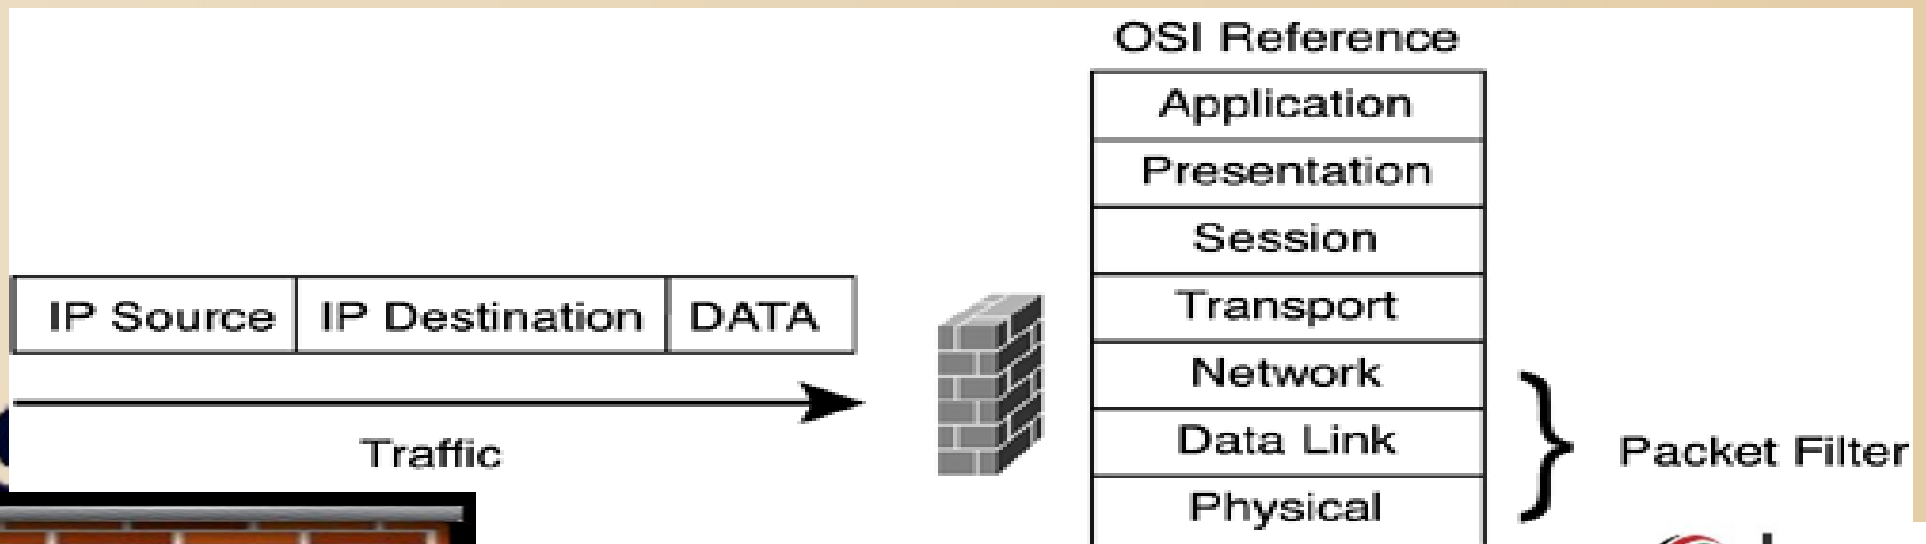
Types of Firewalls

- Firewall are capable of monitoring traffic using different techniques.
- The three types of inspection methodologies are as follows:
 - i. Packet filtering and stateless filtering
 - ii. Stateful filtering
 - iii. Deep packet layer inspection



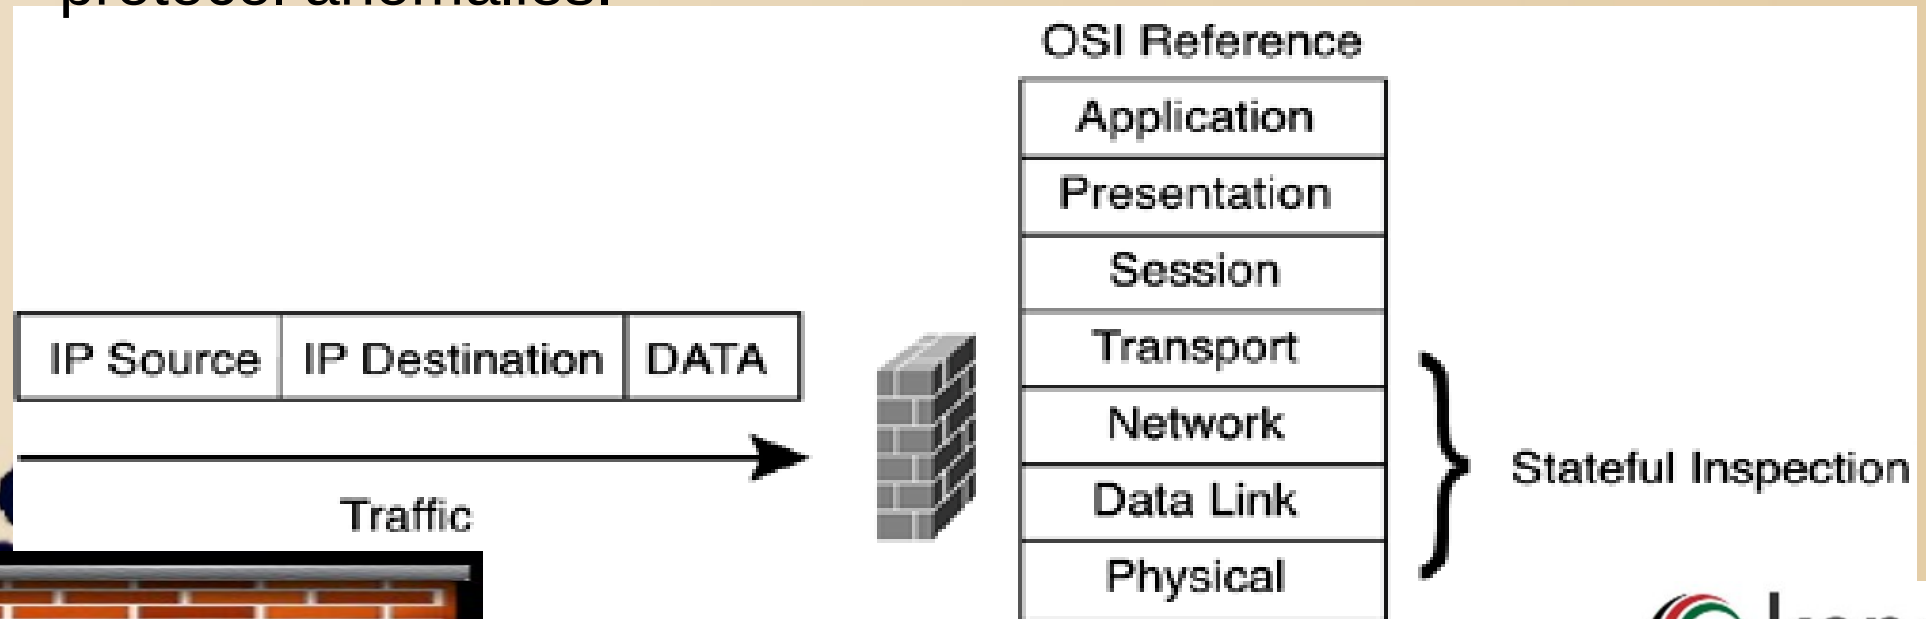
Packet filtering and stateless filtering

- Stateless firewalls are able to inspect source and destination IP addresses and protocol source and destination ports.
- Packets are inspected up to Layer 3 of the OSI model, which is the network layer.



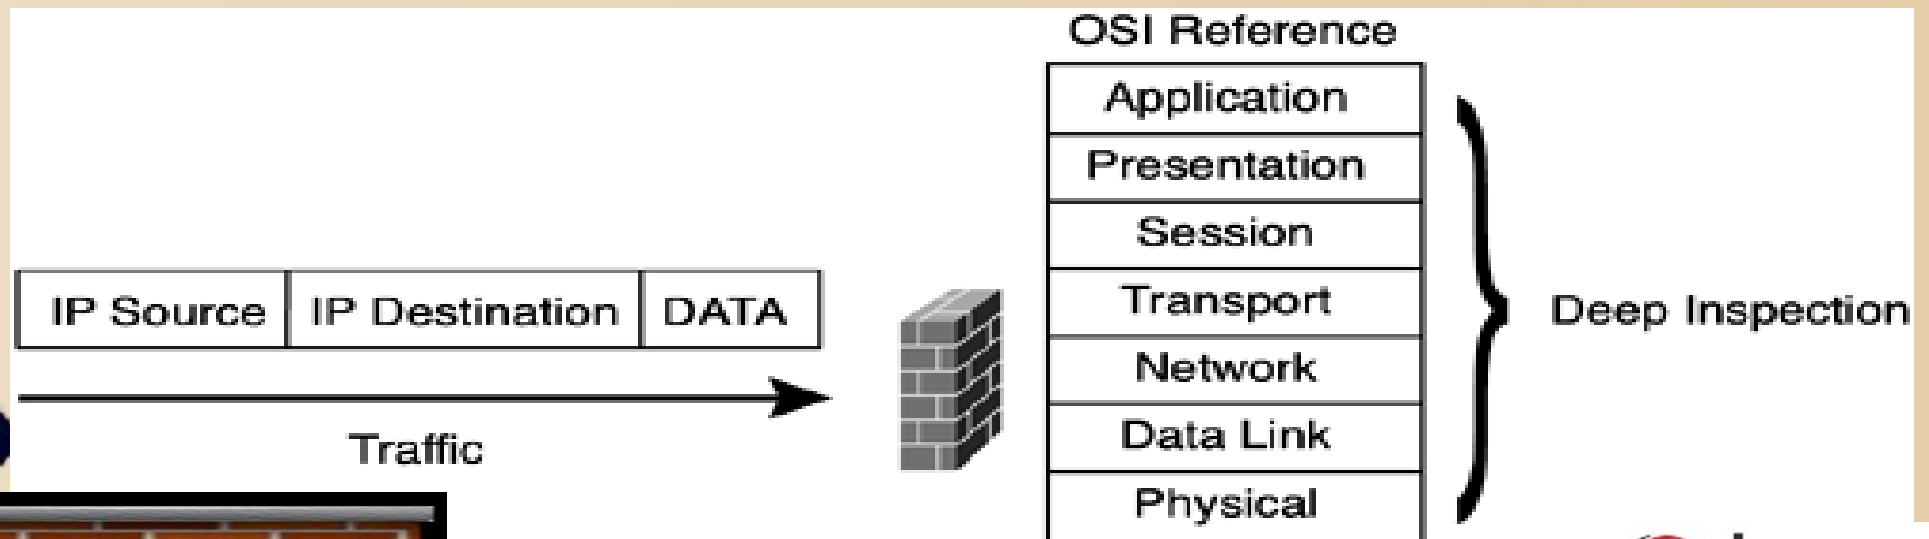
Stateful filtering

- A stateful firewall limits network information from a source to a destination based on the destination IP address, source IP address, source TCP/UDP port, and destination TCP/UDP port.
- Stateful firewalls can also inspect data content and check for protocol anomalies.
- Packets are inspected up to Layer 4 of the OSI model, which is the transport layer. Therefore, stateful firewalls are able to inspect protocol anomalies.



Deep packet layer inspection

- With deep packet layer inspection, the firewall inspects network information from a source to a destination based on the destination IP address, source IP address, source TCP/UDP port, and destination TCP/UDP port.
- It also inspects protocol conformance, checks for application-based attacks, and ensures integrity of the data flow between any TCP/IP devices.



Firewall Packages

- FreeBSD has three different firewall packages built into the base system. They are:
 - IPFILTER (also known as IPF)
 - IPFWALL (also known as IPFW)
 - OpenBSD's PacketFilter (also known as PF).
- FreeBSD also has two built in packages for traffic shaping (basically controlling bandwidth usage): altq and dummynet.
- Dummynet has traditionally been closely tied with IPFW, and ALTQ with PF.
- The reason that FreeBSD has multiple built in packages is that different people

Firewall Packages Cont'd

- A few Linux firewall packages include:
 - Iptables
 - IPCop
 - Shorewall
 - UFW – Uncomplicated Firewall



OpenBSD Packet Filter (PF)



MARTIN NJAU
Systems Administrator
mnjau@kenet.or.ke

28th October 2015



The OpenBSD Packet Filter (PF)

- PF was originally developed by Daniel Hartmeier and is now maintained and developed by the entire OpenBSD team.
- Ported to FreeBSD as of July 2003 and made available in the FreeBSD Ports Collection.
- PF is a complete, full-featured firewall that has optional support for ALTQ (Alternate Queuing). ALTQ provides Quality of Service (QoS) functionality.

Activation

- PF is not enabled by default. To enable it on boot, add the line

`pf_enable="YES"` to the `/etc/rc.conf` file.

- Then run the startup script to load the module:

```
# /etc/rc.d/pf start
```

- Logging support for PF can be enabled by:

- `pflog_enable="YES"`

- Then run the startup script to load the module:

- `# /etc/rc.d/pflog start`

- You can also manually activate and deactivate PF by using the `pfctl` program:

```
# pfctl -e
```

```
# pfctl -d
```

- Note that this just enables or disables PF, it doesn't actually load a ruleset. The ruleset must be loaded separately, either before or after PF is enabled.

Configuration

- PF reads its configuration rules from `/etc/pf.conf` by default at boot time. The sample `pf.conf` can be found in `/usr/share/examples/pf/`.
- The `pf.conf` file has five parts:
 - i) Macros: User-defined variables that can hold IP addresses, interface names, etc.
 - ii) Tables: A structure used to hold lists of IP addresses.
 - iii) Options: Various options to control how PF works.
 - iv) Queueing: Provides bandwidth control and packet prioritization.
 - v) Filter Rules: Allows the selective filtering or blocking of packets as they pass through any of the interfaces. Filter rules can be given parameters to specify network address translation (NAT) and packet redirection.
- Blank lines are ignored, and lines beginning with `#` are treated as comments.

Packet Filtering

- Packet filtering is the selective passing or blocking of data packets as they pass through a network interface. The criteria that pf uses when inspecting packets are based on the Layer 3 (IPv4 and IPv6) and Layer 4 (TCP, UDP, ICMP, and ICMPv6) headers. The most often used criteria are source and destination address, source and destination port, and protocol.
- Filter rules specify the criteria that a packet must match and the resulting action, either block or pass, that is taken when a match is found. Filter rules are evaluated in sequential order, first to last.
- Unless the packet matches a rule containing the quick keyword, the packet will be evaluated against all filter rules before the final action is taken. The last rule to match is the "winner" and will dictate what action to take on the packet.
- There is an implicit pass all at the beginning of a filtering ruleset meaning that if a packet does not match any filter rule the resulting action will be pass.

Rule Syntax

- The general, highly simplified syntax for filter rules is:

```
action [direction] [log] [quick] [on interface] [af] [proto protocol] \  
[from src_addr [port src_port]] [to dst_addr [port dst_port]] \  
[flags tcp_flags] [state]
```

- action - The action to be taken for matching packets, either pass or block.
- direction - The direction the packet is moving on an interface, either in or out.
- log - Specifies that the packet should be logged via pflogd.
- quick - If a packet matches a rule specifying quick, then that rule is considered the last matching rule and the specified action is taken.

Rule Syntax Cont'd

- interface - The name or group of the network interface that the packet is moving through.
- af - The address family of the packet, either inet for IPv4 or inet6 for IPv6.
- protocol - The Layer 4 protocol of the packet.
- src_addr, dst_addr - The source/destination address in the IP header.
- src_port, dst_port - The source/destination port in the Layer 4 packet header.
- tcp_flags - Specifies the flags that must be set in the TCP header when using proto tcp.
- state - Specifies whether state information is kept on packets matching this rule.

Packet Filtering Cont'd

- Default Deny - The recommended practice when setting up a firewall is to take a "default deny" approach.
 - To create a default deny filter policy, the first two filter rules should be:
 - block in all
 - block out all
 - This will block all traffic on all interfaces in either direction from anywhere to anywhere.
 - Passing Traffic - Traffic must now be explicitly passed through the firewall or it will be dropped
- default deny policy.

Macros

▮ #####PF RULES####

- ▮ lan_if="em0" #Your FreeBSD Network Card
- ▮ afnog_v4="196.200.208.0/20" #Afnog's full IPv4 Block
- ▮ afnog_v6="2001:43f8:220::/48" #Afnog's full IPv6 Block
- ▮ bsd_vm_v4="196.200.219.XX" #Your FreeBSD VM IPv4 Block
- ▮ bsd_vm_v6="2001:43f8:220:219:196:200:219:XX" #Your FreeBSD VM's IPv6 Block
- ▮ neighbour_v4="196.200.219.YY" #Your neighbour's IPv4 block
- ▮ neighbour_v6="2001:43f8:220:219:196:200:219:YY" #Your neighbour's IPv6 block

Tables

- ▮ A table is used to hold a group of IPv4 and/or IPv6 addresses. Lookups against a table are very fast and consume less memory and processor time

Tables

```
table <rfc1918> const { 192.168.0.0/16, 172.16.0.0/12,  
10.0.0.0/8 }
```

```
table <kenet-as> const { 41.89.0.0/16, 41.204.160.0/19,  
197.136.0.0/14 }
```

```
table <kenet_stratum2_ntp_servers> const { x.x.x.x/32 }
```

```
table <kenet_intl_links> const { y.y.y.y/24, w.w.w.w/30 }
```

Rules

- ❑ block in all #incoming packets blocked unless
specified in below rules
- ❑ pass out all #whatever communication this server
initiates is allowed
- #Things I want to always allow
- ❑ pass in quick on \$lan_if inet6 proto { tcp, udp, icmp6 } from \$ipv6_ll to
any #allow link local IPv6
- ❑ pass in on \$lan_if from { \$afnog_v4, \$afnog_v6 } to { \$bsd_vm_v4,
\$bsd_vm_v6 }
- ##Add your rules block rules below
- ❑ block in log quick on \$lan_if inet6 proto tcp from \$neighbour_v6 to
\$bsd_vm_v6 port 22
- ❑ block in log quick on \$lan_if inet proto tcp from \$neighbour_v4 to
\$bsd_vm_v4 port { 80, 22 }

Getting Help

- Man pf
- Man pfctl
- ▯ <http://www.openbsd.org/faq/pf/index.html>
- ▯ <http://www.freebsd.org/doc/en/books/handbook/>

GIYF (Google Is Your Friend)



IPTABLES



MARTIN NJAU
Systems Administrator
mnjau@kenet.or.ke

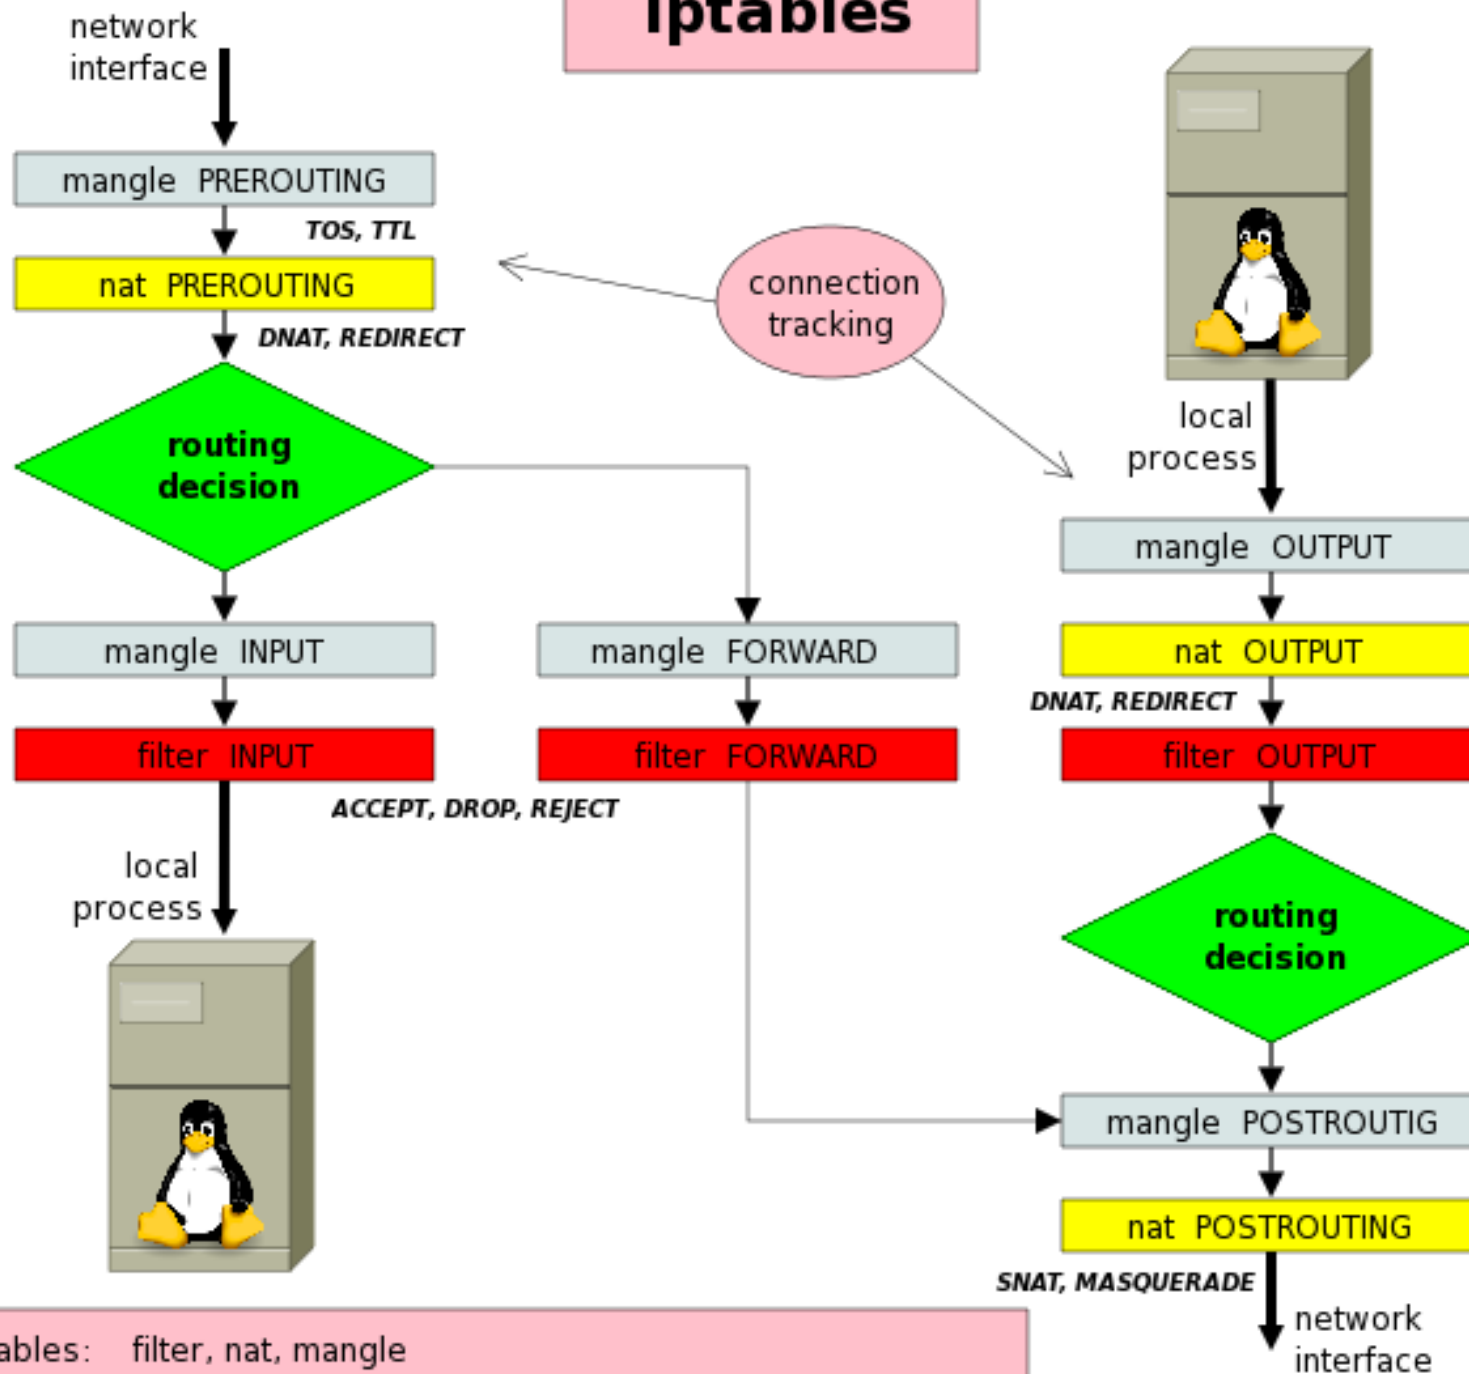
28th October 2015



Introduction

- iptables is the userspace command line program used to configure the Linux 2.4.x and later packet filtering ruleset. It is targeted towards system administrators.
- The iptables package also includes ip6tables. ip6tables is used for configuring the IPv6 packet filter.
- iptables requires a kernel that features the ip_tables packet filter. This includes all 2.4.x and later kernel releases

iptables



tables: filter, nat, mangle

chains: INPUT, OUTPUT, FORWARD, PREROUTING, POSTROUTING

targets: ACCEPT, DROP, REJECT, DNAT, SNAT, MASQUERADE, REDIRECT, LOG, RETURN, TTL, TOS, ...

Basic Commands

- lists your current rules in iptables.

`iptables -L`

- Allowing Established Sessions

`iptables -A INPUT -m state --state ESTABLISHED,RELATED -j ACCEPT`

- Allowing Incoming Traffic on Specific Ports e.g. ssh, http

`iptables -A INPUT -p tcp --dport ssh -j ACCEPT`

`iptables -A INPUT -p tcp --dport 80 -j ACCEPT`

Basic Commands Cont'd

- Blocking Traffic

```
iptables -A INPUT -j DROP
```

- Saving iptables

To save the configuration, you can use iptables-save and iptables-restore.

```
# Iptables-save > /etc/iptables.up.rules
```

Or Create a new rule set:

```
# vi /etc/iptables.up.rules (insert your rules here)
```

```
# vi /etc/network/if-up.d/iptables (insert the following lines:)
```

```
#!/bin/sh
```

```
iptables-restore < /etc/iptables.up.rules
```

Ensure the file is executable (use chmod command)

Example rules

*filter

:INPUT DROP

[0:0]

:FORWARD

Getting Help

- Man iptables
- <http://www.netfilter.org/>

GIYF (Google Is Your Friend)

Q&A.

?



THANK YOU!